

1 This application is submitted in the name of the following
2 inventor(s):

3	<u>Inventor</u>	<u>Citizenship</u>	<u>Residence City and State</u>
4	Henk J. BOTS	United States	Hollister, California

5 The assignee is Network Appliance, Inc., a California
6 corporation having an office at 495 East Java Drive, Sunnyvale, California
7 94089.

8	<u>Title of Invention</u>
---	---------------------------

9 **System To Limit Memory Access When Calculating**
10 **Network Data Checksums**

11 Background of Invention

12 1. *Field of the Invention*

13 This invention relates to the transmission of data through a
14 communications network and incorporates a unique process for calculating
15 data checksums.

16 2. *Related Art*

In computer networks, data is often transmitted from one computer to another using a data transfer protocol. Some of the major protocols in use include Open Systems Interconnection (OSI), Transfer Control Protocol/Internet Protocol (TCP/IP), Decnet, Telnet, XNS, IPX, AppleTalk and others. The invention is intended to work with protocols now in existence or protocols which may be developed in the future, but the

1 invention is particularly useful when the protocol uses a checksum which
2 may be defined to allow transfer of data in small pieces.

3 The transmitting computer obtains the data to be transmitted
4 from a data source (such as a hard disk, DMA or other source) and caches the
5 data in memory. The data is then copied from the application layers to the
6 network layers and ultimately transmitted to the communications network.

7 One problem is that copying the data through a data transfer
8 protocol takes numerous CPU cycles and memory accesses and, thus, slows
9 the transmission process. A method that is sometimes used to avoid copying
10 the data through data transfer protocol is creating a data pointer 114 in the
11 application layer and sending that pointer to the network and transmission
12 layers instead of copying the data itself. Fewer CPU cycles and memory
13 accesses are used to transmit the data pointer 114 than to copy the data
14 itself. The network layers use the data pointer 114 to locate the data cached
15 in memory and transmit it through the communications network.

16 One problem with the method of using a data pointer 114 to
17 speed up the transmission process is that it is that much of the efficiency of
18 transferring the data pointer 114 is lost when the transport layer
19 subsequently needs to read all the data to calculate the checksum. It takes
20 fewer CPU cycles and memory accesses to transfer a data pointer 114 than to
21 copy data through protocol. But CPU cycles and memory accesses are
22 required to read the data cached or stored in memory in order to calculate the
23 checksums. Thus, the economy otherwise gained by transferring a data
24 pointer 114 instead of copying data is diminished because the CPU cycles and
25 memory accesses are required to read the data to calculate the checksums.

Summary of the Invention

The invention provides a method and process for transmitting data using fewer CPU cycles and memory accesses to calculate checksums. The transmitting device obtains data from an internal or external data source and stores that data in memory. The data is then divided into fixed sized zones and checksums are calculated for each zone. The checksums are recorded on a checksum array 115. A data pointer 114 containing an address for the data stored in memory, a description of the data and an address for the checksum array 115 is transferred through data transfer protocol to the network and transmission layers. The network and transmission layers are then able to access and send the data without having to copy or read all of the data to calculate the checksums. This method and process uses fewer CPU cycles and memory accesses to transmit data and is, therefore, more efficient than the prior art.

Brief Description of the Drawings

Figure 1 shows a block diagram of a system for transmitting data across a network.

Figure 2 shows a block diagram of a system for transmitting a data pointer 114 containing the checksum array 115 from an application level through a data transfer protocol to a network level.

Figure 3 shows a block diagram of a system for calculating the data checksum without reading all of the data to be transmitted.

Figure 4 shows a process flow diagram of a preferred method for transmitting data from a data source to a network interface card.

Detailed Description of the Preferred Embodiment

In the following description, preferred and alternative embodiments of the invention are described with regard to process steps and data structures. Embodiments of the invention can be implemented using general-purpose processors or special-purpose processors operating under program control, or other circuits, adapted to particular process steps and data structures described herein. Implementation of the process steps and data structures described herein would not require undue experimentation or further invention.

System Elements

Figure 1 shows a block diagram of a system for transmitting data across a network.

A system 100 includes two or more data transmission devices 110 (such as a computer). In a preferred embodiment, the transmission devices 110 are capable of both transmitting and receiving data at any given time. For the purpose of explanation, when a transmission device is transmitting data, it will be referred to as a “transmitting device” and when it is receiving data, it will be referred to as a “receiving device”. The system also includes a communication network 120, and two or more communication links 121, which link the data transmission devices to the communication network 120.

The data transmission devices 110 include a processor, operational software, data memory 111, and a network interface card 112. The transmission devices 110 are linked either internally or externally to one or more data storage devices (such as a DMA or disk drive) 113.

1 In a preferred embodiment, the communication network 120
2 includes an Internet, intranet, extranet, virtual private network, enterprise
3 network, or another form of communication network. In the preferred
4 embodiment, the communication network 120 is a network capable of routing
5 messages between and among one or more transmission devices. However,
6 there is no particular requirement that the communication network 120
7 comprise an actual network, so long as the communication network 120
8 includes at least some technique for communication between the transmitting
9 and receiving devices 110.

10 The transmitting device 110 obtains data from one or more data
11 storage devices and stores (or “caches”) that data in memory 111. The data is
12 then sent from the memory 111 to the network interface card 112. The data
13 is then sent from the network interface card 112 to the communication
14 network 120 via the communication link.

15 The data travels from the communication network to the
16 network interface card of the receiving device 110 via a communication link
17 121. The data is then sent from the network interface card to the memory
18 111 and ultimately stored on one or more storage device(s) 113.

19 Figure 2 shows a block diagram of a system for transmitting a
20 data pointer 114 containing an address for the data, a description of the data
21 and an address for the checksum array 115 from an application layer
22 through data transfer protocol to network and transmission layers of the
23 transmitting device 110.

24 Data transfer protocols describe how information from a
25 software application in one computer moves through a network medium to a
26 software application in another computer. Some of the major protocols in use
27 include Open Systems Interconnection (OSI), Transfer Control

1 Protocol/Internet Protocol (TCP/IP), Decnet, Telnet, XNS, IPX, AppleTalk
2 and others. The invention is also intended to work with data transfer
3 protocols now in existence or protocols which may be developed in the future,
4 but the invention is particularly useful when the protocol uses a checksum
5 which may be defined to allow transfer of data in small pieces.

6 In the OSI model, for example, information being transferred
7 from a software application in one computer system to a software application
8 in another must pass through each of the seven (7) OSI layers. For the
9 purposes of this application, the seven (7) layers can be apportioned between
10 those layers having more to do with the application and those layers having
11 more to do with the network. Other data transfer protocols, such as TCP/IP,
12 are concerned primarily with the lower network layers of the OSI chain.
13 These two general areas are divided by a transmission layer where the data
14 is transmitted to the network communication system via the data
15 communication link 121.

16 As information is transferred from a data source 113 to memory
17 111 or from one layer to another, the transmission device must check to make
18 sure the integrity of the data remains intact (i.e. all the bytes of data are
19 transferred in the correct order). This is done through the use of a checksum.

20 In the preferred invention, data is obtained from the data source
21 113 and stored (or cached) in memory 111. The data is divided into a series
22 of zones, each containing between 256 and 512 bytes of data. A checksum is
23 calculated for data stored in each of the zones. A checksum array 115
24 containing the checksums for the stored data is created and stored in
25 memory. A data pointer 114 containing an address of where the data is being
26 stored in memory, a description of the data and an address of where the
27 checksum array 115 is being stored in memory is transmitted to the network
28 and transmission layers.

1 Figure 3 shows a block diagram of a system for calculating the
2 data checksum without reading all of the data stored in memory and which is
3 to be transmitted to a receiving device.

4 As stated above, the data held in memory is divided into fixed
5 sized zones, a checksum is calculated for each zone, and an array 115 is
6 created and stored in memory which contains said checksums. During the
7 transmission process, the data transfer protocol divides the data into data
8 chunks and each chunk is transmitted via the transmission layer. The data
9 transfer protocol will not necessarily divide the data according to zones and
10 may split a zone. In order to calculate the checksum for the all the data
11 being sent with each chunk, the data transfer protocol takes the shorter of
12 any divided zone and calculates the checksum for that portion. It then adds
13 that checksum to the checksums for the other zones included within that
14 chunk by accessing the checksum array 115. In this way, the data transfer
15 protocol calculates a checksum for the chunk of data being sent without using
16 the CPU cycles required to read the entire chunk of data.

17 Figure 4 shows a process flow diagram of a preferred method for
18 transmitting data from a data source to a network interface card.

19 At a flow point 410, the transmitting device 110 is coupled to a
20 data source 113, and is ready to transmit data to the receiving device 110 via
21 the communication network 120 and communication link 121.

22 At a step 411, data is transferred from the data source 113 to
23 the memory 111.

24 At a step 412, the data transfer protocol divides the data held in
25 memory into zones of a certain length (preferably 128, 256 or 512 bytes) and

1 calculates a checksum for each zone. The checksums are copied onto a
2 checksum array 115.

3 At a step 413, the transmitting device creates a data pointer 114
4 which includes an address for the data stored in memory, a description of the
5 data stored in memory, and the address of the checksum array 115.

6 At a step 414, the data pointer 114 is transmitted down through
7 data transfer protocol to the network and transmission layers.

8 At a step 415, the data and checksum array 115 are accessed by
9 the network and transmission layers, the data is broken into zones and a
10 checksum is calculated for each chunk to be transmitted using the checksum
11 array 115.

12 At a step 416, the chunk of data to be transmitted is sent from
13 memory to the network interface card.

14 At a step 417, the chunk of data to be transmitted is sent from
15 the network interface card to the communication network 120 via the
16 communications link 121.

17 *Generality of the Invention*

18 The invention has general applicability to various fields of use,
19 not necessarily related to the services described above.

20 Other and further applications of the invention in its most
21 general form, will be clear to those skilled in the art after perusal of this
22 application, and are within the scope and spirit of the invention.

1 *Alternative Embodiments*

2 Although preferred embodiments are disclosed herein, many
3 variations are possible which remain within the concept, scope, and spirit of
4 the invention, and these variations would become clear to those skilled in the
5 art after perusal of this application.